

Modeling Volatility in Prediction Markets

Nikolay Archak
narchak@stern.nyu.edu

Panagiotis G. Ipeirotis
panos@nyu.edu

Department of Information, Operations, and Management Sciences
Leonard N. Stern School of Business, New York University

ABSTRACT

There is significant experimental evidence that prediction markets are efficient mechanisms for aggregating information and are more accurate in forecasting events than traditional forecasting methods, such as polls. Interpretation of prediction market prices as probabilities has been extensively studied in the literature. However there is little research on the volatility of prediction market prices. Given that volatility is fundamental in estimating significance of price movements, it is important to have a better understanding of the volatility of the contract prices.

This paper presents a model of a prediction market with binary payoff on a competitive event involving two parties. In our model, each party has a latent underlying “ability” process that describes its ability to win and evolves as an Ito diffusion. We show that, if the prediction market for this event is efficient and unbiased, the price of the corresponding contract also follows a diffusion and its instantaneous volatility is a function of the current contract price and its time to expiration. In the experimental section, we validate our model on a set of InTrade prediction markets and show that our model is consistent with the observed volatility of contract returns. Our model also outperforms existing volatility models in predicting future contract volatility from historical price data. To demonstrate the practical value of our model, we apply it to pricing options on prediction market contracts, such as those recently introduced by InTrade. Other potential applications of this model include detection of significant market moves and improving forecast standard errors.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

Keywords

prediction market, volatility, forecasting, binary option

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’09, July 6–10, 2009, Stanford, California, USA.

Copyright 2009 ACM 978-1-60558-458-4/09/07 ...\$5.00.

1. INTRODUCTION

There is significant evidence of efficiency and ex-post predictive accuracy in certain types of prediction markets, such as markets for presidential elections [23]. Berg et al. [3] show that the Iowa Electronic Markets significantly outperform polls in predicting the results of national elections. Moreover, they found “*no obvious biases in the market forecasts and, on average, considerable accuracy, especially for large, U.S. election markets*”. Leigh and Wolfers [14] provide statistical evidence that Australian betting markets for the 2004 Australian elections were at least weakly efficient¹ and responded very quickly to major campaign news. Luckner et al. [15] report that prediction markets for the FIFA World Cup outperform predictions based on the FIFA world ranking. According to press releases, Hollywood Stock Exchange prediction market consistently shows 80% accuracy for predicting Oscar nominations [13].

The success of public prediction markets as information aggregation mechanisms led to internal corporate applications of prediction markets for forecasting purposes and as decision support systems [4]. Chen and Plott [7] show that prediction markets on sales forecasting inside HP performed significantly better than traditional corporate forecasting methods in most of the cases. Google has launched internal prediction markets in April 2005; Cowgill et al. [9] report that at Google there were a number of biases such as optimism and overpricing of favorites, however “*as market participants gained experience over the course of our sample period, the biases became less pronounced*”. It is even hypothesized that prediction markets can be used for analysis and evaluation of governmental policies [24].

Interpretation of prediction market prices as probabilities has been extensively studied in the theoretical literature [17, 25]. Nevertheless, little attention so far has been paid to understanding the *volatility* of prediction market prices; a surprising fact, given that volatility is one of the most crucial concepts in the analysis of markets. Volatility has intrinsic interest to prediction market researchers, not only as a measure of market dynamics, but also for its numerous practical applications. Even a simple task of distinguishing “normal” market moves from major events can significantly benefit from having a volatility model.

Example 1 Assume that a company has a binary prediction market contract, asking whether a new product will be ready for launch by December 31st. On September 1st, the price contract increases from 0.50 to 0.55, an increase of 10%.

¹Future returns cannot be predicted from historical prices

The market continues to evolve and on December 15th, the contract price goes down, from 0.95 to 0.855, a decrease of 10%. Are these price movements an indication of an important event? Or are they simply noise?

Given examples like the above, our research question naturally emerges: “If the price of a prediction market contract is the expectation of the actual probability of the event happening, what can we say about the volatility of the contract price?”

To answer this question, we need a model of evolution of the underlying event. Consider a contract that pays \$1 at time T , if and only if some event A happens. If, the outcome of the event is predetermined at time $t < T$, and this information is known to informed (marginal [12]) traders, the price of the contract will not fluctuate in the future. Therefore, for price changes to occur, the event A must either not be predetermined at time $t < T$, or the information about the event must be revealed gradually to all market players.

We model the uncertainty about the event by introducing the notion of “abilities” for the event participants. The “event participants” are the entities that determine whether the event will happen or not. For instance, in 2008 Presidential elections Barack Obama and John McCain were the “event participants”. The abilities of the participants are evolving over time as stochastic processes; therefore their current state reveals only partial information about the future. Furthermore, the larger the time to expiration, the less certain we are about the final state of the process. At the expiration, the state of the “ability” processes defines what is the outcome of the event: the party with the highest “ability” wins.

For our modeling purposes, we assume that abilities evolve as Ito diffusions, a generalized form of a Brownian motion. Ito diffusions are general enough to capture a wide range of behaviors but, at the same time, convenient to work with for deriving analytic results. As the main theoretical contribution of our paper, we show that, under certain assumptions, parameters of the underlying stochastic processes affect the contract price but, given the contract price, *do not* affect its volatility. Moreover, we show that, if we adopt the diffusion model and the underlying “ability” processes are homoscedastic (i.e., the volatility of the ability process does not change over time), the instantaneous volatility of the contract price is fully defined by its current price and the time to expiration.

The rest of the paper is organized as follows. Section 2 gives a short overview of the volatility concept. Section 3 presents our model for pricing of bets in “ideal” prediction markets. Section 4 presents our experimental results obtained for a collection of InTrade prediction markets. Section 5 discusses the experimental results and directions for further research on this topic. Section 6 presents an application of our model to pricing options on prediction market contracts. Finally, Section 7 concludes the paper with a short summary of the theoretical and empirical results.

2. VOLATILITY IN FINANCIAL MARKETS

Volatility is a natural measure of risk in financial markets as it describes the level of uncertainty about future asset returns. Empirical studies of volatility can be traced as far as Mandelbrot [16] who observed that large absolute changes in the price of an asset are often followed by other large absolute changes (not necessarily of the same sign), and small absolute changes are often followed by small absolute changes. This

famous fact is referred to as volatility clustering or volatility persistence and is nowadays a “must have” requirement for any volatility model in the financial literature. It is not until two decades later that the first successful model of volatility forecasting was developed by Engle [11]. The insight of Engle’s ARCH² model was that, in order to capture “volatility clustering,” one should model volatility conditional on previous returns: if the square of the previous return is large one would expect the square of the current return to be large as well. (Note that the *sign* of the return is difficult to predict.) That gave rise to the famous pair of equations:

$$r_t = h_t \varepsilon_t, \quad h_t^2 = \alpha + \beta r_{t-1}^2,$$

where r_t represents return, h_t volatility and ε_t are i.i.d. residuals. Engle’s model was later generalized by Bollerslev [6] (as GARCH³) to allow for lagged volatility in the volatility equation and more advanced generalizations followed in the next two decades [2]. The research on volatility modeling was primarily guided by the observed properties of the empirical distribution of returns in financial markets such as volatility clustering, mean reversion, asymmetry and heavy tails [10].

Our approach differs in that we consider a *binary* prediction market claim not as an equity but as a derivative: a binary option on a couple of latent ability processes. The behavior of a derivative differs in fundamental ways from the behavior of an equity. First, the price is bounded to be between 0.0 and 1.0, unlike an equity that has prices that fluctuate from 0 to infinity. Second, the contract expires at a given time, and the price afterwards is either 0.0 or 1.0, unlike a stock that does not have an expiration date. As we demonstrate later, our model works better than volatility models for equities. Furthermore, our experimental results indicate how to use our model, together with GARCH models in complementary ways, combining the strengths of the two approaches.

3. MODEL

We now introduce our model of volatility. First we describe the general idea of modeling prediction markets using evolving “abilities” and then present the analytic formulations and results.

3.1 Model Setup

A diffusion model can be most naturally introduced for bets on competitive events such as presidential elections or the Super Bowl finals. Consider a prediction market for a simple event in which two parties (McCain vs. Obama or Patriots vs. Giants) compete with each other. The prediction market contract pays \$1 if the first party wins at the expiration of the contract at T , and \$0 otherwise.

Assume that each party has some potential to win. Denote potential of the first party as S_1 , and potential of the second party as S_2 . We consider S_1 and S_2 to be stochastic processes, evolving over time, and refer to them as *ability* processes $S_1(t)$ and $S_2(t)$. Adopting a diffusion approach makes our model similar to a recently published result showing that a simple diffusion model provides a good fit of the evolution of the winning probability in the 2004 Presidential elections market at InTrade.com [8]. However, that paper considers a very restrictive parametric specification (no drift, constant volatility) and does not analyze the volatility of prices.

²AutoRegressive Conditional Heteroscedasticity

³Generalized ARCH




Figure 1: An example of a competition, where the abilities of two parties evolve as Ito Diffusions (top). The bottom figure shows the price of the corresponding prediction market contract. (See Example 2.)

In our model, we assume that, at the expiration T , the strongest party always wins. More formally, the probability of the first party winning, given the abilities $S_1(T)$ and $S_2(T)$ is:

$$\Pr(\text{winner} = 1 | S_1(T), S_2(T)) = \mathbb{I}_+(S_1(T) - S_2(T)) \quad (1)$$

where \mathbb{I}_+ is the indicator function for positive numbers.⁴ (Alternative specifications include smoothing using either logit or probit functions, however, due to space limitations, we do not analyze them in this paper.) Assuming that the prediction market is efficient and unbiased, then the price $\pi(t)$ of the contract at time t , is:

$$\pi(t) = \mathbf{E}_t [\Pr(S_1(T) > S_2(T))] \quad (2)$$

where \mathbf{E}_t represents the expectation taken with respect to all available information at time t . Note that we need only *marginal traders* to share this common information set, as it is the marginal traders who determine the price in a prediction market [12]. Prior study of a political stock market shows that marginal traders are significantly less biased than average traders and efficiently react to news [12], thus providing justification for $\pi(t) = \mathbf{E}_t [\pi(T)]$ if one thinks of \mathbf{E}_t as an expectation of a marginal trader.

In this paper, we model S_1 and S_2 as Ito processes, a generalized form of a Brownian motion. Specifically, we have:

$$dS_i = a_i(S_i, t)dt + b_i(S_i, t)dW_i, \quad i = 1, 2 \quad (3)$$

where $a_i(s, t)$ are *drifts* and $b_i(s, t)$ are *volatilities* of the underlying ability processes, potentially different for each process S_1, S_2 . The processes are driven by two standard Brownian motions W_i that can be correlated with $\text{corr}(W_1, W_2) = \rho_{12}$ ($|\rho_{12}| < 1$).

Example 2 Consider the case of two competing parties, in Figure 1. Party 1 has an ability $S_1(t)$ (red line) with positive drift $\mu_1 = 0.2$ and volatility $\sigma_1 = 0.3$, and party 2 (blue line) has an ability $S_2(t)$ with a negative drift $\mu_2 = -0.2$ and higher volatility $\sigma_2 = 0.6$. Assuming no correlation of abilities, the difference $S(t) = S_1(t) - S_2(t)$ (green line) is a diffusion with

⁴Note that a tie is a zero-probability event.

drift $\mu = 0.4$ and volatility $\sigma = 0.67$. In the bottom plot, you can see the price $\pi(t)$ of the contract, as time evolves.

As shown in the example, the red line (party 1) is for the most time above the blue line (party 2), which causes the green line (the difference) to be above 0. As the contract gets close to expiration, its price gets closer and closer to 1 (i.e., party 1 will win). Close to the end, the blue line catches up, which causes the prediction market contract to have a big swing from almost 1 to 0.5, but then swings back up as party 1 finally finishes at the expiration above party 2.

So far, our model depends on knowing the parameters of the underlying “ability processes.” We now proceed to show that we largely do not need to know the details of the underlying ability processes for our volatility modeling.

3.2 Constant Coefficients

First, we present a relatively simple case, where the drifts and the volatilities of the “ability” processes remain constant over time. Consider a constant coefficients model: $a_i(S_i, t) \equiv \mu_i$, $b_i(S_i, t) \equiv \sigma_i$. In this case, we can consider the difference process ($S = S_1 - S_2$) which can be written as

$$dS = \mu dt + \sigma dW,$$

where $\mu = \mu_1 - \mu_2$, $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2 - 2\rho_{12}\sigma_1\sigma_2}$ and $W = \frac{1}{\sigma}(\sigma_1 W_1 - \sigma_2 W_2)$ is a standard Brownian motion.

By Markovity of our stochastic processes, the price $\pi(t)$ of a prediction market is:

$$\pi(t) = P\{S(T) > 0 | S(t)\},$$

what is equivalent to

$$\pi(t) = P\left\{W(T) - W(t) > -\frac{S(t) + \mu(T-t)}{\sigma}\right\} \quad (4)$$

As W is a Brownian motion, $W(T) - W(t)$ is a normal random variable with mean zero and volatility $\sigma\sqrt{T-t}$, so

$$\pi(t) = \mathcal{N}\left(\frac{S(t) + \mu(T-t)}{\sigma\sqrt{T-t}}\right), \quad (5)$$

where \mathcal{N} is the cumulative density function of the standard normal distribution. While this is a closed-form result, it depends on the unobserved value $S(t)$ and therefore is not directly useful.⁵ We can obtain deeper insight by analyzing the *evolution* of the price process. By applying Ito’s formula [19] to $\pi(t)$, we get:

$$d\pi(t) = \frac{\partial\pi}{\partial t}dt + \frac{\partial\pi}{\partial S}dS + \frac{1}{2}\frac{\partial^2\pi}{\partial S^2}(dS)^2. \quad (6)$$

Now, from Equation 5

$$\begin{aligned} \frac{\partial\pi}{\partial t} &= \frac{1}{2\sigma(T-t)}\phi(\mathcal{N}^{-1}(\pi(t)))\left(\mathcal{N}^{-1}(\pi(t)) - 2\mu\sqrt{T-t}\right), \\ \frac{\partial\pi}{\partial S} &= \frac{1}{\sigma\sqrt{T-t}}\phi(\mathcal{N}^{-1}(\pi(t))), \\ \frac{\partial^2\pi}{\partial S^2} &= -\frac{1}{\sigma^3(T-t)}\phi(\mathcal{N}^{-1}(\pi(t)))\mathcal{N}^{-1}(\pi(t)), \end{aligned}$$

⁵Our derivation is based on the assumption that abilities of each party are public information at time t . Though it is reasonable to assume that this information is well-known to market players, it might not be available to researchers estimating the model (or might be perceived as subjective).

where ϕ stands for probability density function of the standard normal distribution. We can substitute these expressions to the diffusion Equation 6 together with $dS = \mu dt + \sigma dW$ and $(dS)^2 = \sigma^2 dt$. The terms for the drift cancel⁶ and we get:

$$d\pi(t) = 0 dt + \frac{1}{\sqrt{T-t}} \phi(\mathcal{N}^{-1}(\pi(t))) dW.$$

We can see that the instantaneous volatility of $\pi(t)$ (call it $\Sigma(t)$) is given by the expression.

$$\Sigma(t) = \frac{1}{\sqrt{T-t}} \phi(\mathcal{N}^{-1}(\pi(t))), \quad (7)$$

which conditionally depends *only* on the current price $\pi(t)$ and the time to expiration $T-t$ and does *not* depend on any parameters of the underlying latent stochastic processes, given $\pi(t)$ and $T-t$.

3.3 General Case for Binary Markets

In this section we relax the assumption that drifts or volatilities of “ability” processes are constant, though we still have to assume that the coefficients are non-stochastic and that, if the dependency on the current process state is present, it is linear. Note that this theorem covers both the case of a standard Brownian Motion with drift as well as the case of a log-normal Brownian Motion with drift as the latter can be written as $dS = \mu S dt + \sigma S dW$. It also implicitly covers the case of a “threshold” bet that pays \$1 if $S_1(T) > K$ where K is a fixed constant - just take $\mu_2 \equiv \sigma_2 \equiv 0$.

Theorem 1 (Contract pricing in a simple prediction market with two competing parties) Consider a complete probability space (Ω, \mathcal{F}, P) on which we have a two-dimensional Brownian motion $(W_1, W_2; \mathcal{F}_t)$; $0 \leq t \leq T$, where each W_i is a standard Brownian motion, $\text{corr}(W_1(t), W_2(t)) = \rho_{12}$, $|\rho_{12}| < 1$ and \mathcal{F}_t is the natural filtration for $W_{1,2}$. Consider a prediction market for a competitive event such that underlying ability processes S_1 and S_2 are \mathcal{F}_t -measurable and satisfy the diffusion equation

$$dS_i = \mu_i(t)(\alpha S_i + \beta) dt + \sigma_i(t)(\alpha S_i + \beta) dW_i,$$

where $\mu_i(t)$ is a continuous function, $\sigma_i(t)$ is a continuous non-negative function, $\alpha \geq 0$ and $P\{\alpha S_i + \beta > 0\} = 1$. Define the contract price process $\pi(t)$ as

$$\pi(t) = \mathbf{E}_t[\mathbb{I}(S_1(T) > S_2(T))].$$

Under conditions described above, $\pi(t)$ is an Ito’s diffusion with zero drift and instantaneous volatility:

$$\Sigma(t) = \frac{\sigma(t)}{\sqrt{\int_t^T \sigma(u)^2 du}} \phi(\mathcal{N}^{-1}(\pi(t))),$$

where

$$\sigma(s) = \sqrt{\sigma_1^2(s) + \sigma_2^2(s) - 2\rho_{12}\sigma_1(s)\sigma_2(s)}.$$

In other words, $d\pi(t) = \Sigma(t)dW$ where W is some standard Brownian motion with respect to filtration \mathcal{F}_t .

Corollary 1 If the volatilities of the underlying ability processes are constant ($\sigma_i(t) \equiv \sigma_i$):

$$\Sigma(t) = \frac{1}{\sqrt{T-t}} \phi(\mathcal{N}^{-1}(\pi(t))).$$

⁶As they should because of the law of iterated expectations.




Figure 2: Claim Volatility As Function of Its Current Price and Time (claim expires at $T = 1$, current time $t = 0$)

Theorem 1 extends our previous result that, given the current price of the contract and its time to expiration, we can determine its instantaneous volatility without knowing parameters of the underlying ability processes such as their drifts. However, if there are “seasonal” effects in the volatility of the ability processes, our constant coefficient estimate (Equation 7) needs to be scaled by the ratio of the current volatility of the ability processes to the future average volatility of the ability processes until expiration.

For example, if the current contract price is 0.5, the time to expiration is 10 time units and we assume no seasonal effects, then its instantaneous volatility (with respect to the same time units) should be $\frac{1}{\sqrt{10}} \frac{1}{\sqrt{2\pi}} \approx 0.126$. Note that our formula predicts that, given the contract price, volatility is proportional to the inverse square root of the time until expiration, while, given the time to expiration, volatility of the price is a strictly decreasing function of the distance between the contract price and 0.5 and it goes to zero as the contract price approaches 0.0 or 1.0. The dependency of the claim volatility on the contract price and the time to expiration is shown in Figure 2, which demonstrates the “volatility surface” as a function of the contract price and the time to expiration.

Alternatively, we can examine the behavior of price volatility if we do not force the contract price to be constant but let the claim evolve “naturally.” Two interesting questions are:

1. What is the expected instantaneous claim volatility at some future moment of time r ?
2. What is the expected average volatility of the claim from the current moment of time until the future moment of time r ?

The first of these questions is answered by Theorem 2 which says that our best forecast of the instantaneous volatility in the future is the current claim volatility weighted by “seasonal” effects if necessary. The second question is answered by Theorem 3.

Theorem 2 (Instantaneous volatility is a martingale) In the setting of Theorem 1, $\frac{\Sigma(t)}{\sigma(t)}$ is a martingale i.e.

$$\forall r \in [t, T] \quad \mathbf{E}_t \left[\frac{\Sigma(r)}{\sigma(r)} \right] = \frac{\Sigma(t)}{\sigma(t)}.$$

In particular, if we assume that $\sigma_i(t) \equiv \sigma_i$, then $\Sigma(t)$ is a martingale.




Figure 3: Expected contract price conditional on event happening ('x') or not ('o') ($t = 0, T = 1$)

Theorem 3 (Average expected claim volatility) In the setting of Theorem 1, take $r \in (t, T)$ and define

$$\Lambda = \frac{\int_t^r \sigma^2(u) du}{\int_t^T \sigma^2(u) du}$$

i.e. Λ is the volatility-weighted ratio of the time elapsed to the total contract duration (in particular, if we assume that $\sigma_i(t) \equiv \sigma_i$, then $\Lambda = \frac{r-t}{T-t}$). Then,

$$\mathbf{E}_t [(\pi(r) - \pi(t))^2] = \int_0^\Lambda \frac{\phi^2\left(\frac{\mathcal{N}^{-1}(\pi(t))}{\sqrt{1+\lambda}}\right)}{\sqrt{1-\lambda^2}} d\lambda \quad (8)$$

Since $\int_0^1 \frac{\phi^2\left(\frac{\mathcal{N}^{-1}(x)}{\sqrt{1+\lambda}}\right)}{\sqrt{1-\lambda^2}} d\lambda = x - x^2$, it also follows naturally that $\mathbf{E}_t [(\pi(T) - \pi(t))^2] = \pi(t) - \pi(t)^2$, which is the volatility of a Bernoulli trial. This is intuitively expected as, conditional on information set at time t , the price at expiration $\pi(T)$ is a random coin flip, giving 1.0 with probability $\pi(t)$ and 0.0 otherwise.

Corollary 2 (“Ex-post expected” price trajectory) In the setting of Theorem 3,

$$\mathbf{E}_t [\pi(r) | A] = \pi(t) + \frac{1}{\pi(t)} \int_0^\Lambda \frac{\phi^2\left(\frac{\mathcal{N}^{-1}(\pi(t))}{\sqrt{1+\lambda}}\right)}{\sqrt{1-\lambda^2}} d\lambda.$$

Proof: Direct application of Theorem 1 from Pennock et al. [20], which says that, if the prediction market for event A is unbiased, then

$$\mathbf{E} [\pi(t) | \pi(t-1), A] = \pi(t-1) + \frac{\text{Var}\{\pi(t) | \pi(t-1)\}}{\pi(t-1)}. \quad \square$$

Corollary 2 deserves some clarification. Our model was built under assumption that the contract price is a martingale i.e. $\mathbf{E}_t [\pi(r)] = \pi(t)$. It follows that the average of all price trajectories of the claim from point $(t, \pi(t))$ is just a horizontal line. Imagine now that an observer (but not a trader) has access to Oracle that can say whether the event will actually happen or not. Naturally, if the oracle says “yes”, it eliminates all price trajectories converging to zero. Corollary 2 tells us what one would obtain by averaging all remaining trajectories that converge to one. We show the expected price trajectories for several initial price values in Figure 3. Note that all price trajectories converge either to 1.0 (event happens) or to 0.0 (event does not happen).

Variable	Min	Max	Mean	Median	S.Dev.
Num.Obs.	1	1427	375.76	237	334.52
Expiration	2	1933	417.68	347	306.35
Daily dev.	-0.93	0.942	0	0	0.018869
Returns	-0.9975	449	0.0191	0	1.4059
Price	0.001	0.995	0.244	0.075	0.32092

Table 1: Descriptive Statistics for 901 InTrade contracts (338,563 observations)

4. EXPERIMENTAL EVALUATION

In this section, we present the experimental evaluation of our model on a set of InTrade prediction market contracts.

4.1 Data

Our dataset contained daily observations for a collection of InTrade prediction market contracts. Intrade is an online Dublin-based trading exchange web site founded in 2001. The trading unit on InTrade is a contract with a typical settlement value of \$10, which is measured on a 100 points scale. To be consistent with our convention that the winning contract pays \$1, we have renormalized all price data to be in $[0, 1]$ range, so, for example, a 50 points price of InTrade contract is represented by 0.5 in our dataset. The full dataset we analyzed included daily closing price and volume data for a collection of 901 InTrade contracts obtained by periodic crawling of InTrade’s web site.⁷ Table 1 provides the descriptive statistics for our sample. Note that we use term “price difference” or “absolute return” to represent price changes between two consecutive days

$$a_t = p_t - p_{t-1}$$

while just “return” refers to

$$r_t = \frac{p_t - p_{t-1}}{p_{t-1}}$$

A closer examination of the data reveals several interesting facts that must be taken into account in our empirical application. At first, 288,119 of 338,563 observations that we have (i.e., > 85% of the whole dataset) had zero price change since the previous day. Moreover, 287,430 of these observations had zero daily trading volume, what means that most of the time price did not change because of the absence of any trading activity. We presume that absence of activity in many of InTrade prediction markets can be attributed partially to low liquidity of most of the markets in our sample and partially to transaction costs for contracts executed on the exchange. Even if we exclude observations with $a_t = 0$ from our dataset, returns exhibit the following “round number” bias: absolute returns divisible by 5 ticks⁸ occur much more frequently than absolute returns of similar magnitude that are not divisible by 5 ticks. For example, 5 ticks price difference has occurred 2,228 times in our dataset as opposed to 624 times for 4 tick difference and 331 times for 6 tick difference. In fact, more than half of all non-zero price differences in our dataset are divisible by 5 ticks. The bias can be easily seen in Figure 4 that plots absolute return sizes against corresponding frequencies on regular and log scales. Similar bias was also observed for prices.

⁷InTrade keeps historical data for each active contract on the web site, however expired contracts disappear from the web site after certain amount of time.

⁸One InTrade tick is equal to 0.1 InTrade points or \$0.01 in our normalization.




Figure 4: Deviation counts by size in ticks (0.1 InTrade point); zero deviations excluded.

4.2 Model Test

In this section, we assume that the underlying “ability” process is homoscedastic, i.e., $\sigma_i(t) \equiv \sigma_i$ and time t is measured in days, as we use daily returns. Our model suggests that volatility should increase as we approach the expiration of the contract, and that volatility is higher when contract price is close to 0.5 and lower when it approaches 0.0 or 1.0. Specifically, our theoretical model predicts that, for each observation, *conditional on price history*, we have:

$$a_t = h_t \varepsilon_t, \quad (9)$$

$$h_t = \sqrt{\int_0^\Lambda \frac{1}{\sqrt{1-\lambda^2}} \phi^2 \left(\frac{1}{\sqrt{1+\lambda}} \mathcal{N}^{-1}(p_t) \right) d\lambda}, \quad (10)$$

where h_t is the conditional volatility of absolute returns and $\Lambda = \frac{1}{T-t}$. Residuals ε_t are independent identically distributed standard (mean zero, variance one) random variables.⁹ While most of InTrade contracts in our sample have more than two possible outcomes, we believe the qualitative behavior of contract prices can be well captured by our simple result for binary model. To test our statement, we take logs of the absolute value of Equation 9:¹⁰

$$\log(|a_t|) = \mathbf{E}[\log|\varepsilon_t|] + \log(h_t) + \log|\varepsilon_t| - \mathbf{E}[\log|\varepsilon_t|].$$

For notational convenience, let $\gamma = \mathbf{E}[\log|\varepsilon_t|]$ and $\mu_t = \log|\varepsilon_t| - \mathbf{E}[\log|\varepsilon_t|]$. In this case, we have:

$$\log(|a_t|) = \gamma + \log(h_t) + \mu_t,$$

Note that $\mathbf{E}[\mu_t] = 0$ and so we have a regression-like setting even though the residuals are not normally distributed. Therefore, we can use regressions to check the validity of our proposed model. We ran two different tests of our model, taking absolute values of price differences and regressing them on a constant and $\log(h_t)$.

- The first test checks for *presence* of heteroscedasticity effects of time and price as predicted by our theory. The null hypothesis here is that the coefficient on $\log(h_t)$ is equal to zero and we want to reject it.
- The second test checks that the *magnitude* of heteroscedasticity effects is consistent with our theory. The null hypothesis here is that coefficient on $\log(h_t)$ is equal to one and we do not want to reject it.

⁹We do not know what is the distribution of ε_t , however for small values of Λ it must be close to normal.

¹⁰We take logs instead of squares of returns as logs are more robust to outliers in the data. Similar results can be obtained with squares of returns.

Obs.	Var.	Coeff.	Std.Err.	-95%	+95%
all	$\log(h_t)$	0.6298	0.00480	0.6204	0.6392
	constant	-3.3833	0.03084	-3.4437	-3.3229
(0.05, 0.95)	$\log(h_t)$	0.8467	0.01986	0.8078	0.8860
	constant	-2.4653	0.08598	-2.6338	-2.2967
(0.1, 0.9)	$\log(h_t)$	0.9455	0.02830	0.8900	1.0010
	constant	-2.1010	0.11507	-2.3266	-1.8755

Table 2: Testing for presence and magnitude of heteroscedasticity. If volatility is not affected by time to expiration and contract price, then the coefficient of $\log(h_t)$ should be close to 0.0. Coefficients of $\log(h_t)$ close to 1.0 indicate that our model captures well the magnitude of volatility.

Obs.	Var.	Coeff.	Std.Err.	-95%	+95%
all	$\log(h_t)$	0.6438	0.013112	0.61810	0.66951
(0.05, 0.95)	$\log(h_t)$	1.0370	0.039503	0.95957	1.11443
(0.1, 0.9)	$\log(h_t)$	1.1910	0.049832	1.09341	1.28876

Table 3: Three fixed effect regressions of logs of absolute returns

Results of the regression are given in the Table 2 and standard errors are corrected for heteroscedasticity.¹¹ We included zero deviations (but not zero volume trading days), but to avoid taking logs of zero we added a small smoothing factor (10^{-4}) under the $\log(|a_t|)$. In Table 2 we report the results of *three* regressions. The first regression included all observations, the second regression included all observations with price in the (0.05, 0.95) range and the last regression included all observations with price in the (0.1, 0.9) range. We report results with excluded marginal observations because such observations are most seriously affected by the bias depicted in Figure 4. Moreover, there is substantial evidence from the prior research suggesting that people tend to overvalue small probabilities and undervalue near certainties, the so called “favorite-longshot bias” [22]. This effect may be especially strong in our sample as it is skewed towards low price contracts.

As we see from Table 2:

- *Presence of heteroscedasticity*: The null hypothesis for the first test is strongly rejected in all three cases, meaning that there are indeed strong heteroscedasticity effects of the current contract price and the time until contract expiration. In other words, volatility changes over time and across prices.
- *Magnitude of heteroscedasticity*: For the second test, while in the first two cases the null hypothesis is rejected at 95% confidence level, the coefficient on $\log(h_t)$ improves (i.e., approaches the predicted value of 1) as we exclude the marginal observations, which may be affected by the “favorite-longshot bias”. This means that our model captures nicely the magnitude of the volatility, especially for contracts with prices in the (0.1, 0.9) range.

Robustness checks: The results of the pooled regressions might be affected by potential heterogeneity of the contracts. While the basic theory suggests that ε_t are i.i.d. residuals, in

¹¹As one can see from Figure 5, residuals are indeed heteroscedastic: when volatility is high, the residual is high as well.




Figure 5: Sample means of log of absolute returns conditional on fixed price or time to expiration against predictions of our model.

practice we may expect correlations in volatility levels for the same contract, for example, due to different liquidity levels of contracts in the sample. In order to alleviate these concerns we also included contract-specific effects to our regression. In Table 3, we report results of panel data regressions with fixed effects for contracts. Note that the Breusch and Pagan LM test rejects absence of effects ($\chi^2(1) = 5448.76$, $p > \chi^2 = 0.0000$) and the Hausman test does not reject the random effects model ($\chi^2(1) = 0.97$, $p > \chi^2 = 0.3258$).

We also noted that, because of potential auto-correlation in absolute returns, we may underestimate actual standard errors; we tried a random-effects regression with AR(1) disturbances (GLS estimator). Indeed, we found significant auto-correlation in residuals ($\rho = 0.526$); correcting for autocorrelation, though, does not affect the results significantly.

Furthermore, we have replicated all results using the instantaneous volatility from Equation 7 instead of the average daily volatility from Equation 10. As most of the observations in our sample are relatively far from the expiration date, instantaneous volatilities were close to daily averages and the regression results were not significantly different.

As instantaneous volatility expression provides particularly nice separation of time and price effects, we also suggest running a simple visual test in addition to regression based testing. As we have plenty of observations, we can use data to estimate sample means of the $\log(|a_t|)$ conditional on the fixed price or the fixed time to expiration. We can then plot the means against predictions of our model and see if the qualitative behavior is similar. Results that we obtained are presented in Figure 5. Solid lines shown on the picture are plots of $\log(\phi(\mathcal{N}^{-1}(p_{t-1})))$ and $-0.5 \log(T - t)$ shifted by a constant so that they match the data means.

4.3 Volatility Forecasting

So far, in our experimental results we were ignoring potential heteroscedasticity of the latent ability processes and, therefore, potential “volatility clustering” for the contract prices. One can argue that, if volatility clustering is significant, one might be able to forecast claim volatility better by using historical data than by using our theoretical results.

So, we compared our model against models of volatility which use historic volatility to predict future volatility. In our test, we compared three models: GARCH(1,1) [6], our model assuming homoscedasticity of the “ability” processes, and GARCH(1,1) applied to the standardized residuals returned by our model.¹² Note, that our model does not require

¹²The standardized residual is the return r_t normalized by

any historic data, while other two models require estimating parameters of the GARCH process from historical price data.

We performed the comparison on a subsample of 51 In-Trade contracts where each contract represents the Democratic Party Nominee winning Electoral College Votes of a particular state¹³ in the 2008 U.S. Presidential Election, as they contained enough volume and long enough history for training the GARCH models. As we are more interested in forecasting volatility rather than in simply explaining it, we separated observations for each contract into two equal parts: the first part was used to learn parameters of the GARCH processes (on a per contract basis) and the second part was used for evaluation of forecasting accuracy. We have compared all three models in terms of total log-likelihood on the testing part of the sample as this is a natural fit criterion for GARCH models.

In 26 out of 51 cases, Model 3 (GARCH on the standardized residuals from our model) provided the best forecasts and in 17 out of the 25 remaining cases our model outperformed regular GARCH. Overall, the results support the hypothesis that even with significant volatility clustering in the data, our model is generally better in forecasting volatility than the approaches based purely on historical data. Moreover, our model seems to capture effects of time and price on claim volatility that are orthogonal to the heteroscedasticity captured by GARCH. Orthogonality of these two different sources of heteroscedasticity implies that we can significantly improve forecasts of future contract volatility by first normalizing the data using our model and then running GARCH on the standardized residuals, rather than running it directly on the return data.

5. LIMITATIONS

Overall, our experiments show that realized volatility of prediction markets is consistent with what our diffusion-based theory would predict. While match between theory and data is definitely not perfect, the model does provide very good predictions of market volatility without requiring one to know *anything* about the market except for the current price and time to expiration. Why is the model fit not ideal? There are at least three different reasons for that.

The first reason is that our model ignores market micro-structure as well as possible behavioral biases such as “favorite-longshot” and rounding biases. No market (especially prediction market) is ideally liquid - the bid-ask queue always has finite depth and there are usually transaction costs. Even marginal traders are not fully rational as well. It is an interesting research question to examine whether the volatility model can be extended to capture the effects of the structure of the bid-ask queue and/or some standard behavioral biases.

The second reason is that our idea of two underlying “ability” processes being driven by Brownian motions might not represent the real stochastic process driving the event probability. One puts significant restrictions on the contract price dynamics by assuming that they are driven by Brownian motions. Its diffusions result in sample paths that are continuous and short-term price changes that are *almost* normally distributed. However, in our sample we sometimes observe price jumps that are completely improbable if we assume normal

our prediction of the return volatility which is $\frac{h_t}{p_t}$ where h_t is given by Equation 10.

¹³More precisely, one of the fifty states and Washington D.C.




Figure 6: Option prices as a function of the time to option expiration and the strike price. The current underlying contract price is 0.5.

distribution of returns. For example, the contract for Michigan to hold new democratic primaries in 2008 had a 59 cents price drop on March 19th, a move that, according to our estimates, constitutes almost 15 standard deviations. One might want to extend our model to cover cases when the underlying ability processes, in addition to a Brownian motion component, have a Poisson jump component capturing the sudden arrival of new significant information to the market.

The last reason is that we price claims as if traders are risk-neutral. As prediction market bets are somewhat similar to binary options, it is tempting to say that we just borrow risk-neutral pricing approach from the option pricing literature. However, the traditional argument to support risk-neutral pricing of options, i.e., Delta hedging in the underlying instrument does *not* work for prediction markets as the underlying instrument either does not physically exist or cannot be traded. Nevertheless, we can suggest at least three alternative arguments in defense of risk-neutral pricing. At first, in some prediction markets traders may indeed behave as risk-neutral either because the market uses “play money” instead of real ones [21] or because trader’s participation is limited. Next, there is theoretical evidence that, under certain conditions, prediction market prices may be close to the mean population beliefs even if the traders are not risk-neutral [25], so risk-neutral pricing might be a valid approach even in the presence of certain risk-aversion. Finally, as we already described in the beginning of the paper, there is significant experimental evidence that prediction market prices are unbiased estimates of the actual event probability. Note that this last argument is our main justification in this paper. Without trying to answer the question of “why is it so?”, we just adopted risk-neutral pricing to see where the theory leads us. The results we have obtained in the experimental part of this paper, might be seen as a joint test of market efficiency, non-bias assumptions, as well as of our volatility model.

6. APPLICATION: PRICING OPTIONS ON PREDICTION MARKET CONTRACTS

This section presents an application of our model for pricing options on prediction market contracts. Options are popular financial instruments with numerous applications such as risk hedging or speculation on volatility. A classic “vanilla” call option on a security provides the right but not the obligation to buy a specified quantity of the security at a set strike price

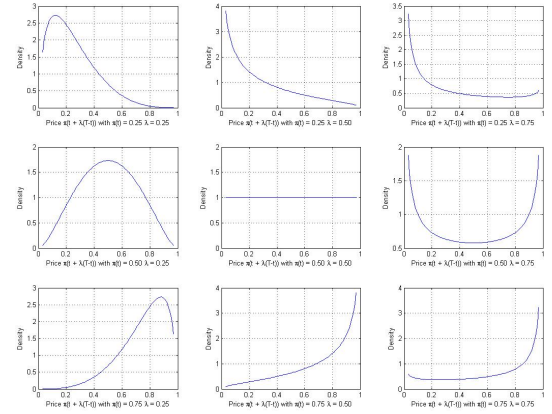


Figure 7: Risk-neutral densities for future contract prices. Rows correspond to $\pi(t) = 0.25, 0.5, 0.75$ and columns correspond to $\lambda = 0.25, 0.5, 0.75$.

at a certain expiration date.¹⁴ A major breakthrough in option pricing was achieved by Fischer Black and Myron Scholes who obtained a closed form solution for pricing options, now known as the Black-Scholes formula [5, 18].

In this section, we consider binary options on prediction market claims, like those recently introduced by InTrade. Such binary option will pay \$1 if on option’s expiration date T' , the underlying contract price is larger than the strike price (K). The underlying contract expires at time T , with $T' \leq T$.

For example, InTrade’s option *X.15OCT.OBAMA*. > 74.0 pays 100 points (\$1 in our normalization) if on October 15, 2008 (T'), the price of the 2008.PRES.OBAMA contract is higher than 74 points (K). The expiration date for the underlying contract is on November 9, 2008 (T).

As we already assume risk-neutrality, it is easy to define the option price:

$$c(t) = P\{\pi(T') > K | \pi(t)\}. \quad (11)$$

We know the evolution process for the underlying contract, so we only need to calculate its expectation. The result is stated by the following theorem.

Theorem 4 (Pricing options on prediction market contracts in a simple prediction market with two competing parties) *In the setting of Theorem 1, take $T' \in (t, T)$ and consider a binary option on the contract $\pi(s)$ with payoff*

$$c(t) = P\{\pi(T') > K | \pi(t)\}.$$

Define $\lambda = \frac{\int_t^{T'} \sigma^2(u) du}{\int_t^T \sigma^2(u) du}$ i.e. λ is the volatility-weighted ratio of the time to the option expiration to the time to the contract expiration. (In particular, if we assume that $\sigma_i(t) \equiv \sigma_i$, then $\lambda = \frac{T' - t}{T - t}$.) Then the option price is given by:

$$c(t) = \mathcal{N}\left(\sqrt{\frac{1}{\lambda}} \mathcal{N}^{-1}(\pi(t)) - \sqrt{\frac{1}{\lambda}} - 1 \mathcal{N}^{-1}(K)\right). \quad (12)$$

Several interesting observations can be made from Equation 12. At first, the option price is strictly increasing in the

¹⁴In this section, we consider European style options only. An American style option will give owner the right to exercise it on or before the expiration date.

current contract price and strictly decreasing in the strike price. Moreover, as time to the option expiration converges to time to the contract expiration ($\lambda \rightarrow 1$), the option price converges to the current contract price ($c(t) \rightarrow \pi(t)$). Finally, if the current contract price is 0.5 and the option expires halfway to the contract expiration (as weighted by volatilities of the “ability” processes), then the option price is equal to the strike price $c_t = K$. To give the reader better intuition of how the option price behaves with varying parameters λ and K we include a plot of option prices for a fixed value of the current price ($\pi(t) = 0.5$) in Figure 6.

One can also take derivative of the option price with respect to the strike price to retrieve the risk-neutral density of the contract price at time T' . The result is

$$f(p) = \frac{\phi\left(\frac{\sqrt{\frac{1}{\lambda}}\mathcal{N}^{-1}(\pi(t)) - \sqrt{\frac{1}{\lambda} - 1}\mathcal{N}^{-1}(p)}{\phi(\mathcal{N}^{-1}(p))}\right)\sqrt{\frac{1}{\lambda} - 1}}{\phi(\mathcal{N}^{-1}(p))} \quad (13)$$

We plotted the result for different values of the current price and λ in Figure 7.

7. IMPLICATIONS AND CONCLUSIONS

Although volatility is one of the most widely studied concepts in financial markets, little attention, so far, has been given to understanding volatility of prediction markets. This paper is the first attempt to provide a theoretical model of prediction market volatility. In doing so, we assume unbiased and efficient prediction market and assume that the event being predicted is driven by a pair of latent diffusion processes.¹⁵ Combination of both assumptions generates some interesting theoretical results, like the instantaneous volatility of a contract on a binary event depending only on the current contract price and the time to expiration.

The volatility results we obtained bear certain similarity to the family of ARCH models [11]. The main difference is that ARCH models represent conditional volatility of returns in stock market as a function of previous returns (the well-known effect of volatility clustering), while our model suggests that conditional volatility of absolute returns in prediction market is a function of the current price and the time to expiration. The second difference is that, while ARCH models are usually empirical and rely on past data, our model of conditional heteroscedasticity can be derived theoretically from a stochastic model of latent ability processes.

While our theory is based on a model of an “ideal” market, our experimental results for a collection of 901 real InTrade prediction markets show that volatility patterns of real prediction markets are consistent with what our model predicts, especially if we exclude marginal (very high or very low priced) observations from the dataset. Our results for a sample 51 of InTrade contracts on 2008 Presidential Elections demonstrate that our model is better in forecasting volatility than GARCH applied to historical price data, although the best performance is obtained by combining both models. Further practical applications of our results may include detection of significant market moves, improving forecast standard errors in prediction markets, pricing conditional prediction markets and pricing options on prediction markets.

Despite the limitations (outlined in Section 5), we hope that this study provides solid foundation on which future work on prediction market volatility can build.

¹⁵Our results can be extended to events with more than two participants. Refer to Appendix of [1] for details.

8. ACKNOWLEDGMENTS

The authors would like to thank the participants at the TUILES seminar at New York University, Jason Ruspini, and David Pennock for their helpful comments on earlier versions of this paper.

9. REFERENCES

- [1] N. Archak and P. G. Ipeirotis. Modeling volatility in prediction markets, Sept. 2008. Working Paper.
- [2] A. Bera and M. Higgins. ARCH models: Properties, estimation and testing. *Journal of Economic Surveys*, 7(4):305–366, 1993.
- [3] J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research, 2003. Working Draft, The Handbook of Experimental Economics Results.
- [4] J. Berg and T. Rietz. Prediction markets as decision support systems. *Information Systems Frontiers*, 5(1):79–93, 2003.
- [5] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [6] T. Bollerslev. Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics*, 31:307–327, 1986.
- [7] K.-Y. Chen and C. Plott. Information aggregation mechanisms: Concept, design and field implementation for a sales forecasting problem, 2002.
- [8] M. K. Chen, J. Jonathan E. Ingersoll, and E. H. Kaplan. Modeling a presidential prediction market. *Management Science*, 54(8):1381–1394, Aug. 2008.
- [9] B. Cowgill, J. Wolfers, and E. Zitzewitz. Using prediction markets to track information flows: Evidence from Google, 2008. Working Paper.
- [10] R. Engle and A. Patton. What good is a volatility model? *Quantitative Finance*, 1(2):237–245, Feb. 2001.
- [11] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1008, July 1982.
- [12] R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *American Economic Review*, 82(5):1142–61, December 1992.
- [13] HSX. Hollywood stock exchange traders hit 80% of oscar nominations for the 80th annual academy awards. Available at <http://www.hsx.com/about/press/080123.htm>, 2008.
- [14] A. Leigh and J. Wolfers. Competing approaches to forecasting elections: Economic models, opinion polling and prediction markets. *The Economic Record*, 82(258):325–340, 2006.
- [15] S. Luckner, J. Schroder, and C. Slamka. On the forecast accuracy of sports prediction markets. In *Negotiation, Auctions, and Market Engineering*, volume 2, pages 227–234. Springer Berlin Heidelberg, 2008.
- [16] B. Mandelbrot. The variation of certain speculative prices. *Journal of Business*, 36(4):394–419, 1963.
- [17] C. F. Manski. Interpreting the predictions of prediction markets. *Economic Letters*, 91(3):425–429, June 2006.
- [18] R. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1):141–183, 1973.
- [19] B. Oksendal. *Stochastic Differential Equations: An Introduction with Applications (Universitext)*. Springer, December 2005.
- [20] D. M. Pennock, S. Debnath, E. J. Glover, and C. L. Giles. Modeling information incorporation in markets with application to detecting and explaining events. In *18th Conference on Uncertainty in Artificial Intelligence (UAI 2002)*, pages 405–413, 2002.
- [21] E. Servan-Schreiber, J. Wolfers, D. M. Pennock, and B. Galebach. Prediction markets: Does money matter? *Electronic Markets*, 14(3):243–251, 2004.
- [22] R. Thaler and W. Ziemba. Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2):161–174, 1988.
- [23] J. Wolfers and E. Zitzewitz. Prediction markets, 2004. Working Paper.
- [24] J. Wolfers and E. Zitzewitz. Using markets to evaluate policy: The case of the iraq war, 2004. Working Paper.
- [25] J. Wolfers and E. Zitzewitz. Interpreting prediction market prices as probabilities, 2006. Working Paper.

APPENDIX

Proof of Theorem 1

Define the function $f(S) = \int_{x_0}^S \frac{1}{\alpha x + \beta} dx$, where x_0 is some fixed value. Let $Y_i(t) = f(S_i(t))$ be a new stochastic process. Note that f is strictly increasing on the domain of S , so we have that

$P\{Y_1(T) > Y_2(T)\} = P\{S_1(T) > S_2(T)\}$ i.e. instead of original processes we can consider transformed ones. After applying Ito's formula in integral representation we obtain

$$Y_i(T) = Y_i(t) + \int_t^T \left(\mu_i(u) - \frac{\alpha}{2} \sigma_i^2(u) \right) du + \int_t^T \sigma_i(u) dW_i(u).$$

So $Y_1(T) > Y_2(T)$ if and only if

$$X(t) = \int_t^T \sigma_1(u) dW_1(u) - \int_t^T \sigma_2(u) dW_2(u)$$

is greater than $Y(t) =$

$$Y_2(t) - Y_1(t) + \int_t^T \left(\mu_2(t) - \mu_1(t) + \frac{\alpha}{2} (\sigma_2^2(t) - \sigma_1^2(t)) \right) dt.$$

By properties of Brownian motion, $X(t)$ is a normal variable with mean zero and variance

$$\sigma_{t,T}^2 = \int_t^T \sigma(u)^2 du,$$

where

$$\sigma(t) = \sqrt{\sigma_1^2(t) + \sigma_2^2(t) - 2\rho_{12}\sigma_1(t)\sigma_2(t)}.$$

It follows that $\pi(t) = \mathcal{N}\left(\frac{Y(t)}{\sigma_{t,T}}\right)$. By applying Ito's formula to $\pi(t)$, we can calculate its drift and volatility. Note, however, that the drift must be zero by the law of iterated expectations. The volatility term is just

$$\frac{\partial \pi}{\partial Y_1} \sigma_1(t) dW_1 + \frac{\partial \pi}{\partial Y_2} \sigma_2(t) dW_2 = \frac{\sigma(t)}{\sigma_{t,T}} \phi(\mathcal{N}^{-1}(\pi(t))) dW,$$

where

$$W(t) = \frac{\sigma_1(t)W_1(t) - \sigma_2(t)W_2(t)}{\sigma(t)}$$

is a standard Brownian motion. That proves the Theorem. To prove the Corollary just note that if $\sigma(u) \equiv c$ on $[t, T]$, then $\sigma_{t,T} = c\sqrt{(T-t)}$. \square

Proof of Theorem 2

We will use notation from the proof of Theorem 1, in particular, definitions of $\sigma(t)$ and $\sigma_{t,T}$. Let $V(t) = \Sigma(t) \frac{\sigma_{t,T}}{\sigma(t)} = \phi(\mathcal{N}^{-1}(\pi(t)))$. This can be written as $V(t) = U(\pi(t))$, where $U(x) = \phi(\mathcal{N}^{-1}(x))$. Note that $U'(x) = -\mathcal{N}^{-1}(x)$ and $U''(x) = -\frac{1}{U(x)}$. We are now ready to apply Ito's formula in integral form which gives us:

$$V(s) = V(t) - \int_t^s \mathcal{N}^{-1}(\pi(u)) d\pi(u) - \int_t^s \frac{(d\pi(u))^2}{2U(\pi(u))}. \quad (14)$$

Note, that $\int_t^s \mathcal{N}^{-1}(\pi(u)) d\pi(u)$ is an Ito integral and its integrand is square integrable (follows from the construction of the "ability" processes), so its expectation is zero. The second integral is

$$\int_t^s \frac{(d\pi(u))^2}{2U(\pi(u))} = \int_t^s \frac{U^2(\pi(u))\sigma^2(u)du}{2\sigma_{u,T}^2 U(\pi(u))} = \int_t^s \frac{\sigma^2(u)V(u)du}{2\sigma_{u,T}^2}.$$

After applying conditional expectation to Equation 14 we get

$$\mathbf{E}_t[V(s)] = V(t) - \frac{1}{2} \int_t^s \frac{\sigma^2(u)}{\sigma_{u,T}^2} \mathbf{E}_t[V(u)] du,$$

here we used Fubini's theorem to put the expectation operator under the integral. Now, if we define $f(u) = \mathbf{E}_t[V(u)]$, we have an integral equation

$$f(s) = f(t) - \frac{1}{2} \int_t^s \frac{\sigma^2(u)}{\sigma_{u,T}^2} f(u) du.$$

After taking the derivative of both sides with respect to s and rearranging terms we get:

$$\frac{\partial}{\partial s} (\log f(s)) = \frac{1}{2} \frac{\partial}{\partial s} \left(\log \left(\sigma_{s,T}^2 \right) \right).$$

Using the fact that $\frac{\partial}{\partial s} \log \left(\sigma_{s,T}^2 \right) = -\frac{\sigma^2(s)}{\sigma_{s,T}^2}$, one can see that all solutions of this equation are of the form $f(s) = C\sigma_{s,T}$. But then,

$$\mathbf{E}_t[\Sigma(s)] = \frac{\sigma(s)}{\sigma_{s,T}} \mathbf{E}_t[V(s)] = \frac{\sigma(s)f(s)}{\sigma_{s,T}} \equiv C\sigma(s). \quad \square$$

Proof of Theorem 3

We will use notation from the proof of Theorem 1, in particular, definitions of $\sigma(t)$ and $\sigma_{t,T}$. First, note that

$$\mathbf{E}_t(\pi(r) - \pi(t))^2 = \mathbf{E}_t \left(\int_t^r \Sigma(u) dW \right)^2 = \int_t^r \mathbf{E}_t[\Sigma(u)^2] du, \quad (15)$$

where the last step is obtained by applying Ito's isometry [19]. Now, the easiest proof is obtained by noting that the evolution process of the prediction market contract

$$d\pi(t) = \frac{\sigma(t)}{\sigma_{t,T}} \phi(\mathcal{N}^{-1}(\pi(t))) dW,$$

does not depend on parameters of the "ability" processes except for σ . As we will obtain the same result with any valid values of μ, α and β , we can as well take the simplest possible set: $\mu \equiv 0, \alpha \equiv 0, \beta \equiv 1$. With these settings $Y(t) = Y_1(t) - Y_2(t)$ is just a Brownian motion and $\pi(u) = \mathcal{N}\left(\frac{Y(u)}{\sigma_{u,T}}\right)$ and $\Sigma(u) = \frac{\sigma(u)}{\sigma_{u,T}} \phi\left(\frac{Y(u)}{\sigma_{u,T}}\right)$. By using normality of increments of a standard Brownian motion:

$$\mathbf{E}_t[\Sigma(u)^2] = \int_{-\infty}^{\infty} \frac{1}{\sigma_{t,u}} \phi\left(\frac{\delta}{\sigma_{t,u}}\right) \left[\frac{\sigma(u)}{\sigma_{u,T}} \phi\left(\frac{Y(t) + \delta}{\sigma_{u,T}}\right) \right]^2 d\delta.$$

Simple algebraic calculations give $\phi\left(\frac{\delta}{\sigma_{t,u}}\right) \phi^2\left(\frac{Y(t) + \delta}{\sigma_{u,T}}\right) =$

$$\phi\left(\delta \frac{\sqrt{2\sigma_{t,u}^2 + \sigma_{u,T}^2}}{\sigma_{t,u}\sigma_{u,T}} + C\right) \phi^2\left(\frac{Y(t)}{\sqrt{2\sigma_{t,u}^2 + \sigma_{u,T}^2}}\right).$$

Only the first term depends on δ , so after integration:

$$\mathbf{E}_t[\Sigma(u)^2] = \frac{\sigma^2(u)/\sigma_{t,T}^2}{\sqrt{1-\lambda(u)}\sqrt{1+\lambda(u)}} \phi^2\left(\frac{\mathcal{N}^{-1}(\pi(t))}{\sqrt{1+\lambda(u)}}\right),$$

where $\lambda(u) = \frac{\sigma_{t,u}^2}{\sigma_{t,T}^2}$. Note that $d(\lambda(u)) = \frac{\sigma(u)^2}{\sigma_{t,T}^2}$. If we substitute this expression to Equation 15 and change the variable as $u \mapsto \lambda$ we immediately obtain Equation 3. \square

Proof of Theorem 4

We will use notation from the proof of Theorem 1. One can note that the evolution process of the prediction market contract $d\pi(t) = \Sigma(t)dW$ does not depend on parameters of the "ability" processes, so we will obtain the same result with any valid values of μ, α and β . Take $\mu \equiv 0, \alpha \equiv 0, \beta \equiv 1$. With these settings $Y(t) = Y_1(t) - Y_2(t)$ is just a Brownian motion and $\pi(t) = \mathcal{N}\left(\frac{Y(t)}{\sigma_{t,T}}\right)$, so $P\{\pi(T') > K|\pi(t)\}$ can be written as

$$P\left\{ \mathcal{N}\left(\frac{Y(T')}{\sigma_{T',T}}\right) > K \mid \mathcal{N}\left(\frac{Y(t)}{\sigma_{t,T}}\right) = \pi(t) \right\}.$$

The resulting formula can be obtained by simple algebraic manipulations and the fact that $Y(T') - Y(t)$ is normally distributed with mean zero and standard deviation $\sigma_{t,T'}$. \square